arXiv:1702.03828v1 [math.OC] 13 Feb 2017

SHARPNESS, RESTART AND ACCELERATION

VINCENT ROULET AND ALEXANDRE D'ASPREMONT

ABSTRACT. The Łojasievicz inequality shows that sharpness bounds on the minimum of convex optimization problems hold almost generically. Here, we show that sharpness directly controls the performance of restart schemes. The constants quantifying sharpness are of course unobservable, but we show that optimal restart strategies are fairly robust, and searching for the best scheme only increases the complexity by a logarithmic factor compared to the optimal bound. Overall then, restart schemes generically accelerate accelerated methods.

1. INTRODUCTION

We focus on solving an unconstrained convex optimization problem written

$$\text{minimize } f(x) \tag{P}$$

where f is a convex function defined on $\mathbb{R}^n$. The complexity of solving this problem using first order methods is controlled by smoothness assumptions on the gradient of f. Here, we will assume that ∇f is Hölder continuous, i.e. that there are constants $s \geq 1$ and $L > 0$ such that

$$\|\nabla f(x) - \nabla f(y)\| \leq L\|x - y\|^{s-1}, \quad \text{for every } x, y \in \mathbb{R}^n, \tag{Smooth}$$

where $\|.\|$ is the Euclidean norm. This condition generalizes the typical smoothness condition. The Universal Fast Gradient method [Nesterov, 2015] solves problem (P) with optimal complexity, converging to an ϵ accurate solution in

$$O\left(\epsilon^{-\frac{2}{3s-2}}\right) \text{ iterations.}$$

Here we make the following additional assumption on the *sharpness* of minimizers

$$\mu d(x, X^*)^r \leq f(x) - f^*, \quad \text{for every } x \in K \tag{Sharp}$$

where f^* is the minimum of f, $K \subset \mathbb{R}^n$ is a compact set, $d(x, X^*)$ the Euclidean distance from x to the set $X^* \subset K$ of minimizers of f, and $r > 0$, $\mu > 0$ are constants. This defines a *lower bound* on the function, which generalizes similar lower bounds produced by strong or uniform convexity conditions for example. Strong and uniform convexity guarantee linear [Nesterov, 2003] convergence and faster polynomial rates [Iouditski and Nesterov, 2014] respectively when solving (P). Here, we use restart schemes to show that the sharpness also leads to faster rates.

Real analytic functions all satisfy (Sharp) locally, and this inequality is then known as Łojasievicz's inequality [Łojasiewicz, 1963]. This result has recently been generalized to a much wider class of non-smooth convex functions [Łojasiewicz, 1993, Bolte et al., 2007]. Since (Sharp) essentially measures the sharpness of minimizers, it holds somewhat generically, meaning we should typically expect better convergence rates using restart schemes when r or μ are nontrivial. On the other hand this inequality is purely descriptive, linking sharpness with faster rates, as we have no hope of ever observing either r or μ, and deriving adaptive schemes will be key for practical performance.

Restart schemes were already studied for strongly or uniformly convex functions [Nemirovskii and Nesterov, 1985, Nesterov, 2007, Iouditski and Nesterov, 2014]. In particular, [Nemirovskii and Nesterov, 1985]

Date: February 14, 2017.
2010 *Mathematics Subject Classification.* 90C25, 90C06.
Key words and phrases. Acceleration, Restart, Łojasiewicz inequality.

link a "strict minimum" condition akin to (Sharp) with faster convergence rates using restart schemes which form the basis of our results, but do not study the cost of adaptation. More recently, several heuristics [O'Donoghue and Candes, 2015, Su et al., 2014, Giselsson and Boyd, 2014] studied adaptive restart schemes to speed up convergence of optimal methods. Sharpness was also used to prove linear converge in matrix games in [Gilpin et al., 2012].

Our contribution here is to derive optimal scheduled restart schemes for convex functions with Hölder continuous gradient and to show that these schemes can be made adaptive with nearly optimal complexity (up to a squared log term). We also show that when the optimal value of problem (P) (or an accurate termination criterion) is known, then convergence can be further accelerated by stopping inner iterations once a certain target precision has been reached. Here too we show that restart schemes can be made adaptive to the Łojasievicz's inequality parameters, with a much smaller cost than in the general case.

The paper is organized as follows. In Section 2, we recall some key results regarding Łojasievicz's inequality on convex subanalytic functions, and study the link between sharpness and smoothness. In Section 3, we present scheduled restart strategies for smooth convex functions satisfying the Łojasievicz inequality, together with adaptive variants. In Section 4, we generalize these results to functions with Hölder continuous gradients. In Section 5, we present adaptive restart schemes when f^* is known. Finally, we describe numerical experiments in Section 6.

Notations. In what follows, we write f^* the minimum of f, and X^* the set of minimizers, while $d(x, S)$ is the Euclidean distance from a point $x \in \mathbb{R}^n$ to a set $S \subset \mathbb{R}^n$. $\lceil x \rceil$ and $\lfloor x \rfloor$ are respectively the smallest integer larger than or equal to x and the largest integer smaller than or equal to x.

2. Sharpness

The notion of sharp minimum has been extensively studied in [Łojasiewicz, 1963, Łojasiewicz, 1993, Burke et al., 2002]. We briefly recall some consequences of these results below. Further details and a simple proof of the main inequality in dimension one are given in Appendix A.1.

2.1. The Łojasevicz Inequality.

The Łojasewicz inequality [Łojasiewicz, 1963, Łojasiewicz, 1993] shows that if a function $f : \mathbf{dom}\, f \subset \mathbb{R}^n \to \mathbb{R}$ is analytic then for every compact set $K \subset \mathbf{dom}\, f$ there exists $r > 0$ and $\mu > 0$ such that

$$\mu d(x, X^*)^r \leq f(x) - f^*, \quad \text{for every } x \in K.$$

An equivalent and better known formulation shows that if x^* is a critical point of f, i.e. $\nabla f(x^*) = 0$, then there is $\theta \in [0, 1[$ such that

$$\frac{|f(x) - f(x^*)|^\theta}{\|\nabla f(x)\|}$$

is bounded in a neighborhood of x^*. Compared with the previous formulation, $\theta = 1 - 1/r$ here. For $\theta = 1/2$, this matches the gradient dominated property introduced by [Polyak, 1963] and recently used by e.g. [Karimi et al., 2016] to simplify linear convergence proofs for a number of classical algorithms. This also directly implies that any local minimizer is in fact a global minimizer, even when f is nonconvex. These inequalities are easy to prove for univariate functions which is illustrated in appendix A.1, but they are much harder to prove in arbitrary dimension n and we refer the reader to e.g. [Bierstone and Milman, 1988] for details.

The earlier references cited above require f to be analytic around its critical points. However, Łojasewicz's inequality can be generalized to a much broader class of functions, including certain C^1 functions [Kurdyka, 1998]. In particular, Bolte et al. [2007] show an extension of Łojasewicz's inequality holding for a broad class of nonsmooth functions called subanalytic. We recall their definition in Appendix A.2. Inequality (Sharp) is shown using topological arguments that are far from constructive, hence outside of some particular cases (e.g. strong convexity), we cannot assume that the constants in (Sharp) are known, even

approximately. Nevertheless, (Sharp) has a fairly intuitive interpretation as a *sharpness measure* for the minimum of f and can easily be linked to the smoothness exponent as we detail below.

2.2. **Sharpness and smoothness .** Let f be a convex function on $\mathbb{R}^n$ whose gradient is Hölder continuous, i.e. satisfies (Smooth) with parameters (s, L). This property ensures that,

$$f(x) \leq f(y) + \nabla f(y)^T(x-y) + \frac{L}{s}\|x-y\|^s, \quad \text{for every } x, y \in \mathbb{R}^n.$$

For a given $x \in \mathbb{R}^n$, setting y to be the projection of x onto X^*, this yields the following *upper bound* on suboptimality

$$f(x) - f^* \leq \frac{L}{s} d(x, X^*)^s. \tag{1}$$

Now assume that f satisfies the Łojasiewicz inequality (Sharp) on a set K with parameters (r, μ). Combining (1) and (Sharp) we get for every $x \in K$,

$$\frac{s\mu}{L} \leq d(x, X^*)^{s-r}.$$

This means that we necessarily have $s \leq r$ if we take $x \to X^*$. Moreover if $s < r$, this last inequality can only be valid on a bounded set, *i.e.* either smoothness or sharpness or both are valid only on a bounded set. In the following, we write

$$\kappa \triangleq L^{\frac{2}{s}}/\mu^{\frac{2}{r}} \tag{2}$$

a generalized condition number for the function f. If $r = s = 2$, this matches the classical condition number of the function. Finally, we write

$$\tau \triangleq 1 - \frac{s}{r}, \tag{3}$$

another condition number based on the ratio of powers in inequalities (Smooth) and (Sharp).

3. Scheduled restarts for smooth convex problems

In this section we seek to solve (P) assuming that the function f is smooth, i.e. satisfies (Smooth) with $s = 2$ and $L > 0$. Without further assumptions on f, the optimal algorithm to solve the smooth convex optimization problem (P) is Nesterov's Accelerated Gradient Method [Nesterov, 1983]. Given an initial point x_0, this algorithm outputs, after t iterations, a point $y = \mathcal{A}(x_0, t)$ such that

$$f(y) - f^* \leq \frac{cL}{t^2} d(x_0, X^*)^2, \tag{4}$$

where $c > 0$ denotes a universal constant (whose value will be allowed to vary in what follows, with $c = 4$ here). More details about Nesterov's algorithm are given in Appendix B.2.

In what follows, we will also assume that f satisfies (Sharp) with parameters (r, μ) on a set $K \supseteq X^*$, which means

$$\mu d(x, X^*)^r \leq f(x) - f^*, \quad \text{for every } x \in K.$$

As mentioned before if $r > s = 2$, this property is necessarily local, *i.e.* K is bounded. We assume then that given a starting point $x_0 \in \mathbb{R}^n$, sharpness is satisfied on the sublevel set $\{x,\ f(x) \leq f(x_0)\}$. Remark that if this property is valid on an open set $K \supset X^*$, it will also be valid on any compact set $K' \supset K$ with the same exponent r but a potentially lower constant μ. Scheduled restarts that we present rely on a global sharpness hypothesis on the sublevel set defined by the initial point and are not locally adaptive to constant μ. Restarts on criterion that we present in section 5 with the additional hypothesis that f^* is known, show for their part local adaptivity. We now describe a restart scheme exploiting this extra regularity assumption to improve the computational complexity of solving problem (P) using accelerated methods.

3.1. **Scheduled restarts.** Here we schedule the number of iterations t_k made by Nesterov's algorithm before the k^{th} restart. Our scheme is described below.

Algorithm 1 Scheduled restarts for smooth convex minimisation **(RESTART)**

Inputs: $x_0 \in \mathbb{R}^n$ and a sequence t_k for $k = 1, \dots, R$.
for $k = 1, \dots, R$ **do**

$$x_k := \mathcal{A}(x_{k-1}, t_k) \qquad \text{(RESTART)}$$

end for
Output: $\hat{x} := x_R$

In the results that follow, the sequence t_k is chosen a priori to ensure linear convergence. We begin by showing a technical lemma which assumes linear convergence holds, and connects the precision reached, the total number of inner iterations N, and the growth of t_k.

Lemma 3.1. *Suppose setting*

$$t_k = Ce^{\alpha k}, \quad k = 1, \dots, R$$

for some $C > 0$ and $\alpha \geq 0$ in the (RESTART) *scheme generates a sequence of points x_k whose objective values $f(x_k)$ converge linearly, i.e.*

$$f(x_k) - f^* \leq \nu e^{-\gamma k}, \tag{5}$$

for all $k \geq 0$ with $\nu \geq 0$ and $\gamma \geq 0$. Writing $N = \sum_{k=1}^{R} t_k$ the total number of iterations, the precision reached at the last point $\hat{x}$ generated by (RESTART) *is then given by,*

$$f(\hat{x}) - f^* \leq \nu \exp(-\gamma N/C), \quad \textit{when } \alpha = 0,$$

and

$$f(\hat{x}) - f^* \leq \frac{\nu}{(\alpha e^{-\alpha} C^{-1} N + 1)^{\frac{\gamma}{\alpha}}}, \quad \textit{when } \alpha > 0.$$

Proof. When $\alpha = 0$, $N = RC$, and inserting this in (5) at the last point $\hat{x}$ where $k = R$ yields the desired result. On the other hand, when $\alpha > 0$, we have $N = \sum_{k=1}^{R} t_k = Ce^{\alpha} \frac{e^{\alpha R} - 1}{e^{\alpha} - 1}$, which gives

$$R = \frac{\log\left(\frac{e^{\alpha}-1}{e^{\alpha} C} N + 1\right)}{\alpha}.$$

Inserting this in (5) at the last point, we get

$$f(\hat{x}) - f^* \leq \nu \exp\left(-\frac{\gamma}{\alpha} \log\left(\frac{e^{\alpha} - 1}{e^{\alpha} C} N + 1\right)\right) \leq \frac{\nu}{(\alpha e^{-\alpha} C^{-1} N + 1)^{\frac{\gamma}{\alpha}}}$$

where we used $e^x - 1 \geq x$. This yields the second part of the result. ■

The last approximation in the case $\tau > 0$ simplifies the analysis that follows without significantly affecting the bounds. Also, in practice t_k is an integer, but Lemma C.1 in the Appendix shows that using $\tilde{t}_k = \lceil t_k \rceil$ does not significantly affect the bounds above.

We now analyze restart schedules t_k that ensure linear convergence. Our choice of t_k will heavily depend on the ratio between r and s (with $s = 2$ for smooth functions here), incorporated in the parameter $\tau = 1 - s/r$ defined in (3). Below, we show that if $\tau = 0$, a constant schedule is sufficient to ensure linear convergence. When $\tau > 0$, we need a geometrically increasing number of iterations for each cycle.

Proposition 3.2. *Let f be a smooth convex function satisfying* (Smooth) *with parameter L, and $x_0 \in \mathbb{R}^n$. Assume that f satisfies the sharpness property* (Sharp) *with parameters (r, μ) on a set K, such that $\{x,\ f(x) \leq f(x_0)\} \subset K$. We write x_k the sequence generated by the* (RESTART) *scheme to solve* (P) *with iteration schedule $t_k = C^*_{\kappa,\tau} e^{\tau k}$, for $k = 1, \dots, R$, where*

$$C^*_{\kappa,\tau} \triangleq e^{1-\tau} (c\kappa)^{\frac{1}{2}} (f(x_0) - f^*)^{-\frac{\tau}{2}}, \tag{6}$$

with κ and τ are defined in (2) and (3) respectively. The precision reached at the last point $\hat{x}$ is given by,

$$f(\hat{x}) - f^* \leq \exp\left(-2e^{-1}(c\kappa)^{-\frac{1}{2}}N\right)(f(x_0) - f^*), \quad \textit{when } \tau = 0, \tag{7}$$

while

$$f(\hat{x}) - f^* \leq \frac{f(x_0) - f^*}{\left(\tau e^{-1}(f(x_0) - f^*)^{\frac{\tau}{2}}(c\kappa)^{-\frac{1}{2}}N + 1\right)^{\frac{2}{\tau}}}, \quad \textit{when } \tau > 0, \tag{8}$$

where $N = \sum_{k=1}^{R} t_k$ is the total number of iterations.

Proof. Our strategy is to choose t_k such that the objective is geometrically decreasing, *i.e.*

$$f(x_k) - f^* \leq e^{-\gamma k}(f(x_0) - f^*), \tag{9}$$

for some $\gamma \geq 0$ depending on the choice of t_k. This directly holds for $k = 0$ and any $\gamma \geq 0$. Combining (Sharp) with the complexity bound in (4), we get

$$f(x_k) - f^* \leq \frac{c\kappa}{t_k^2}(f(x_{k-1}) - f^*)^{\frac{2}{r}}.$$

Assuming that (9) is satisfied at iteration $k-1$ for a given γ, we have

$$f(x_k) - f^* \leq \frac{c\kappa e^{-\gamma\frac{2}{r}(k-1)}}{t_k^2}(f(x_0) - f^*)^{\frac{2}{r}},$$

and to ensure (9) at iteration k, we impose

$$\frac{c\kappa e^{-\gamma\frac{2}{r}(k-1)}}{t_k^2}(f(x_0) - f^*)^{\frac{2}{r}} \leq e^{-\gamma k}(f(x_0) - f^*).$$

Rearranging terms in last inequality, using τ defined in (3), we get

$$t_k \geq e^{\frac{\gamma(1-\tau)}{2}}(c\kappa)^{\frac{1}{2}}(f(x_0) - f^*)^{-\frac{\tau}{2}}e^{\frac{\tau\gamma}{2}k}. \tag{10}$$

For a given $\gamma \geq 0$, choosing $t_k = Ce^{\alpha k}$ where

$$\begin{aligned} C &= e^{\frac{\gamma(1-\tau)}{2}}(c\kappa)^{\frac{1}{2}}(f(x_0) - f^*)^{-\frac{\tau}{2}} \\ \alpha &= \frac{\tau\gamma}{2}, \end{aligned}$$

and Lemma 3.1 then yields

$$f(\hat{x}) - f^* \leq \exp\left(-\gamma e^{-\frac{\gamma}{2}}(c\kappa)^{-\frac{1}{2}}N\right)(f(x_0) - f^*), \quad \text{when } \tau = 0,$$

while

$$f(\hat{x}) - f^* \leq \frac{(f(x_0) - f^*)}{\left(\frac{\tau}{2}\gamma e^{-\frac{\gamma}{2}}(c\kappa)^{-\frac{1}{2}}(f(x_0) - f^*)^{\frac{\tau}{2}}N + 1\right)^{\frac{2}{\tau}}}, \quad \text{when } \tau > 0.$$

These bounds are minimal for $\gamma = 2$, which yields the desired result. ∎

We also get the following corollary on more generic restart schedules, which will prove useful in further results on adaptivity.

Corollary 3.3. *Let f be a smooth convex function satisfying (Smooth) with parameter L, and $x_0 \in \mathbb{R}^n$. Assume that f satisfies the sharpness property (Sharp) with parameters (r, μ) on a set K, such that $\{x,\ f(x) \leq f(x_0)\} \subset K$. We write x_k the sequence generated by the (RESTART) scheme to solve (P). Given general schedules of the form*

$$\begin{cases} t_k = C & \textit{if } \tau = 0, \\ t_k = Ce^{\alpha k} & \textit{if } \tau > 0, \end{cases}$$

we have the following complexity bounds

$$f(\hat{x}) - f^* \leq \left(\frac{c\kappa}{C^2}\right)^{\frac{N}{C}} (f(x_0) - f^*) \quad \text{if } \tau = 0 \text{ and } C \geq C^*_{\kappa,0}, \tag{11}$$

while

$$f(\hat{x}) - f^* \leq \frac{f(x_0) - f^*}{(\alpha e^{-\alpha} C^{-1} N + 1)^{\frac{2}{\tau}}} \quad \text{if } \tau > 0 \text{ and } C \geq C(\alpha), \tag{12}$$

where

$$C(\alpha) \triangleq e^{\frac{\alpha(1-\tau)}{\tau}} (c\kappa)^{\frac{1}{2}} (f(x_0) - f^*)^{-\frac{\tau}{2}}, \tag{13}$$

and $N = \sum_{k=1}^{R} t_k$ is the total number of iterations.

Proof. Given general schedules of the form

$$\begin{cases} t_k = C & \text{if } \tau = 0, \\ t_k = Ce^{\alpha k} & \text{if } \tau > 0, \end{cases}$$

the best value of γ satisfying condition (10) for any $k \geq 0$ in Proposition 3.2 are given by

$$\begin{cases} \gamma = \log\left(\frac{C^2}{c\kappa}\right) & \text{if } \tau = 0 \text{ and } C \geq C^*_{\kappa,0}, \\ \gamma = \frac{2\alpha}{\tau} & \text{if } \tau > 0 \text{ and } C \geq C(\alpha). \end{cases}$$

As above, plugging these values into the bounds of Lemma 3.1 yields the desired result. ■

When $\tau = 0$, bound (7) matches the classical complexity bound for smooth strongly convex functions. When $\tau > 0$ on the other hand, bound (8) highlights a much faster convergence rate than accelerated gradient methods. The sharper the function (*i.e.* the smaller r), the faster the convergence. This matches the lower bounds for optimizing smooth and sharp functions functions [Nemirovskii and Nesterov, 1985, Page 6] up to constant factors. Also, setting $t_k = C^*_{\kappa,\tau} e^{\tau k}$ yields continuous bounds on precision, i.e. when $\tau \to 0$, bound (8) converges to bound (7). This also shows that for τ near zero, constant restart schemes are almost optimal.

3.2. **Adaptive scheduled restart.** The previous restart schedules depend on sharpness parameters (r, μ) in (Sharp). In general, these values are neither observed nor known a priori. Making our restart scheme adaptive is thus crucial to its practical performance. Fortunately, we show below that a simple logarithmic grid search strategy on these parameters is enough to guarantee nearly optimal performance.

We run several schemes with a fixed number of inner iterations N to perform a log-scale grid search on τ and κ. We define these schemes as follows.

$$\begin{cases} \mathcal{S}_{i,0} : \text{(RESTART) scheme with } t_k = C_i, \text{ where } C_i = 2^i, \\ \\ \mathcal{S}_{i,j} : \text{(RESTART) scheme with } t_k = C_i e^{\tau_j k}, \text{ where } C_i = 2^i \text{ and } \tau_j = 2^{-j}. \end{cases} \tag{14}$$

We stop these schemes when the total number of inner algorithm iterations has exceed N, *i.e.* at the smallest R such that $\sum_{k=1}^{R} t_k \geq N$. The size of the grid search in C_i is naturally bounded as we cannot restart the algorithm after more than N total inner iterations, so $i \in [1, \ldots, \lfloor \log_2 N \rfloor]$. We will also show that when τ is smaller than $1/N$, a constant schedule performs as well as the optimal geometrically increasing schedule, which means we can also choose $j \in [1, \ldots, \lceil \log_2 N \rceil]$. The following result details the convergence of this method.

Proposition 3.4. *Let f be a smooth convex function satisfying* (Smooth) *with parameter L, $x_0 \in \mathbb{R}^n$ and N a given number of iterations. Assume that f satisfies the sharpness property* (Sharp) *with parameters (r, μ) on a set K, such that $\{x,\ f(x) \leq f(x_0)\} \subset K$. We run schemes $\mathcal{S}_{i,j}$ defined in* (14) *to solve* (P) *for $i \in [1, \ldots, \lfloor \log_2 N \rfloor]$ and $j \in [0, \ldots, \lceil \log_2 N \rceil]$, stopping each time after N total inner algorithm iterations i.e. for R such that $\sum_{k=1}^{R} t_k \geq N$. Assuming N is large enough, i.e. $N \geq 2C^*_{\kappa,\tau}$, and that, if $\frac{1}{N} > \tau > 0$, $C^*_{\kappa,\tau} > 1$.*

If $\tau = 0$, there exists $i \in [1, \dots, \lfloor \log_2 N \rfloor]$ such that scheme $\mathcal{S}_{i,0}$ achieves a precision given by

$$f(\hat{x}) - f^* \leq \exp\left(-e^{-1}(c\kappa)^{-\frac{1}{2}}N\right)(f(x_0) - f^*).$$

If $\tau > 0$, there exist $i \in [1, \dots, \lfloor \log_2 N \rfloor\}$ and $j \in [1, \dots, \lceil \log_2 N \rceil]$ such that scheme $\mathcal{S}_{i,j}$ achieves a precision given by

$$f(\hat{x}) - f^* \leq \frac{f(x_0) - f^*}{\left(\tau e^{-1}(c\kappa)^{-\frac{1}{2}}(f(x_0) - f^*)^{\frac{\tau}{2}}(N-1)/4 + 1\right)^{\frac{2}{\tau}}}.$$

Overall, running the logarithmic grid search has a complexity $(\log_2 N)^2$ times higher than running N iterations in the optimal scheme.

Proof. Denote $N' = \sum_{k=1}^{R} t_k \geq N$ the number of iterations of a scheme $\mathcal{S}_{i,j}$. We necessarily have $N' \leq 2N$ for our choice of C_i and τ_j. Hence the cost of running all methods is of the order $(\log_2 N)^2$.

If $\tau = 0$ and $N \geq 2C^*_{\kappa,0}$, we have $i = \lceil \log_2 C^*_{\kappa,0} \rceil \leq \lfloor \log_2 N \rfloor$. Therefore $\mathcal{S}_{i,0}$ has been run and we can use bound (11) to show that the last iterate $\hat{x}$ satisfies

$$f(\hat{x}) - f^* \leq \left(\frac{c\kappa}{C_i^2}\right)^{\frac{N}{C_i}}(f(x_0) - f^*).$$

Using that $C^*_{\kappa,0} \leq C_i \leq 2C^*_{\kappa,0}$, we get

$$\begin{aligned} f(\hat{x}) - f^* &\leq \left(\frac{c\kappa}{(C^*_{\kappa,0})^2}\right)^{\frac{N}{2C^*_{\kappa,0}}}(f(x_0) - f^*) \\ &\leq \exp\left(-e^{-1}(c\kappa)^{-\frac{1}{2}}N\right)(f(x_0) - f^*). \end{aligned}$$

If $\tau \geq \frac{1}{N}$ and $N \geq 2C^*_{\kappa,\tau}$, we have $j = \lceil -\log_2 \tau \rceil \leq \lceil \log_2 N \rceil$ and $i = \lceil \log_2 C^*_{\kappa,\tau} \rceil \leq \lfloor \log_2 N \rfloor$. Therefore scheme $\mathcal{S}_{i,j}$ has been run. As $C_i \geq C^*_{\kappa,\tau} \geq C(\tau_j)$, where $C(\tau_j)$ is defined in (13), we can use bound (12) to show that the last iterate $\hat{x}$ of scheme $\mathcal{S}_{i,j}$ satisfies

$$f(\hat{x}) - f^* \leq \frac{f(x_0) - f^*}{\left(\tau_j e^{-\tau_j} C_i^{-1} N + 1\right)^{\frac{2}{\tau}}}.$$

Finally, by definition of i and j, $2\tau_j \geq \tau$ and $C_i \leq 2C^*_{\kappa,\tau}$, so

$$\begin{aligned} f(\hat{x}) - f^* &\leq \frac{f(x_0) - f^*}{\left(\tau e^{-\tau_j}(C^*_{\kappa,\tau})^{-1}N/4 + 1\right)^{\frac{2}{\tau}}} \\ &= \frac{f(x_0) - f^*}{\left(\tau e^{-1}(c\kappa)^{-\frac{1}{2}}(f(x_0) - f^*)^{\frac{\tau}{2}}N/4 + 1\right)^{\frac{2}{\tau}}}, \end{aligned}$$

where we concluded by expanding $C^*_{\kappa,\tau} = e^{1-\tau}(c\kappa)^{\frac{1}{2}}(f(x_0) - f^*)^{-\frac{\tau}{2}}$ and using that $\tau \geq \tau_j$.

If $\frac{1}{N} > \tau > 0$ and $N > 2C^*_{\kappa,\tau}$, we have $i = \lceil \log_2 C^*_{\kappa,\tau} \rceil \leq \lfloor \log_2 N \rfloor$, so scheme $\mathcal{S}_{i,0}$ has been run. Its iterates x_k satisfy, with $1 - \tau = 2/r$,

$$\begin{aligned} f(x_k) - f^* &\leq \frac{c\kappa}{C_i^2}(f(x_{k-1}) - f^*)^{\frac{2}{r}} \\ &\leq \left(\frac{c\kappa}{C_i^2}\right)^{\left(1-(1-\tau)^k\right)/\tau} (f(x_0) - f^*)^{(1-\tau)^k} \\ &\leq \left(\frac{c\kappa(f(x_0) - f^*)^{-\tau}}{C_i^2}\right)^{\left(1-(1-\tau)^k\right)/\tau} (f(x_0) - f^*). \end{aligned}$$

Now $C_i \geq C^*_{\kappa,\tau} = e^{1-\tau}(c\kappa)^{\frac{1}{2}}(f(x_0) - f^*)^{-\frac{\tau}{2}}$ and $C_i R \geq N$, therefore last iterate $\hat{x}$ satisfies

$$f(\hat{x}) - f^* \leq \exp\left(-2(1-\tau)\frac{1-(1-\tau)^{N/C_i}}{\tau}\right)(f(x_0) - f^*).$$

As $N \geq C_i$, since

$$h(\tau) = \frac{(1-\tau)\left(1-(1-\tau)^{\frac{N}{C_i}}\right)}{1-(1-\tau)}$$

is decreasing with τ and $\frac{1}{N} > \tau > 0$, we have

$$\begin{aligned} f(\hat{x}) - f^* &\leq \exp\left(-2(N-1)\left(1-\left(1-\frac{1}{N}\right)^{N/C_i}\right)\right)(f(x_0) - f^*) \\ &\leq \exp\left(-2(N-1)\left(1-\exp\left(-\frac{1}{C_i}\right)\right)\right)(f(x_0) - f^*) \\ &\leq \exp\left(-2\frac{N-1}{C_i}\left(1-\frac{1}{2C_i}\right)\right)(f(x_0) - f^*). \end{aligned}$$

having used the facts that $(1+ax)^{\frac{b}{x}} \leq \exp(ab)$ if $ax \geq -1$, $\frac{b}{x} \geq 0$ and $1 - x + \frac{x^2}{2} \geq \exp(-x)$ when $x \geq 0$. By assumption $C^*_{\kappa,\tau} \geq 1$, so $C_i \geq 1$ and finally

$$\begin{aligned} f(\hat{x}) - f^* &\leq \exp\left(-\frac{N-1}{C_i}\right)(f(x_0) - f^*) \\ &\leq \exp\left(-\frac{N-1}{2C^*_{\kappa,\tau}}\right)(f(x_0) - f^*) \\ &\leq \frac{f(x_0) - f^*}{\left(\tau(C^*_{\kappa,\tau})^{-1}(N-1)/4 + 1\right)^{\frac{2}{\tau}}} \\ &\leq \frac{f(x_0) - f^*}{\left(\tau(f(x_0) - f^*)^{\frac{\tau}{2}} e^{-1}(c\kappa)^{-\frac{1}{2}}(N-1)/4 + 1\right)^{\frac{2}{\tau}}}. \end{aligned}$$

using the fact that $e^\tau \geq 1$. ■

As showed in Corrolary 3.3, scheduled restart schemes are theoretically efficient only if the algorithm itself makes a sufficient number of iterations to decrease the objective value. Therefore we need N large enough to ensure the efficiency of the adaptive method. If $\tau = 0$, we naturally have $C^*_{\kappa,0} \geq 1$, therefore if $\frac{1}{N} > \tau > 0$ and N is large, assuming $C^*_{\kappa,\tau} \approx C^*_{\kappa,0}$, we get $C^*_{\kappa,\tau} \geq 1$. Convergence of this adaptive method is similar to the one of Nesterov [2003] to optimize smooth strongly convex functions in the sense that we lose approximately a log factor of the condition number of the function. However our assumptions

are weaker and we are able to tackle all regimes of the sharpness property, *i.e.* any exponent $r \in [2, +\infty]$. We end this section by analyzing the behavior of gradient descent in light of the sharpness assumption.

3.3. **Comparison to gradient descent.** Given only the smoothness hypothesis, the gradient descent algorithm, recalled in Appendix B.3, starts from a point x_0 and outputs iterates $x_t = \mathcal{G}(x_0, t)$ such that

$$f(x_t) - f^* \leq \frac{L}{t} d(x_0, X^*)^2,$$

While accelerated methods use the last two iterates to compute the next one, simple gradient descent algorithm uses only the last iterate, so the algorithm can be seen as (implicitly) restarting at each iteration. Its convergence can therefore be written for $k \geq 1$,

$$f(x_{k+t}) - f^* \leq \frac{L}{t} d(x_k, X^*)^2. \tag{15}$$

and we analyze it in light of the restart interpretation using the sharpness property in the following theorem.

Theorem 3.5. *Let f be a smooth convex function with parameter L satisfying* (Smooth) *and $x_0 \in \mathbb{R}^n$. Assume that f satisfies the sharpness property* (Sharp) *with parameters (r, μ) on a set K, such that $\{x,\, f(x) \leq f(x_0)\} \subset K$. We write $x_t = \mathcal{G}(x_0, t)$ the iterate sequence generated by the gradient descent algorithm started at x_0 to solve* (P)*. Define*

$$t_k = e^{1-\tau} \kappa (f(x_0) - f^*)^\tau e^{\tau k}.$$

The precision reached after $N = \sum_{k=1}^n t_k$ iterations is given by

$$f(x_N) - f^* \leq \exp\left(-e^{-1}\kappa^{-1} N\right)(f(x_0) - f^*), \quad \text{if } \tau = 0$$

while

$$f(x_N) - f^* \leq \frac{f(x_0) - f^*}{(\tau e^{-1} \kappa^{-1} (f(x_0) - f^*)^\tau N + 1)^{\frac{1}{\tau}}}, \quad \text{if } \tau > 0.$$

Proof. For a given $\gamma \geq 0$, we construct a subsequence $x_{\phi(k)}$ of x_t such that

$$f(x_{\phi(k)}) - f^* \leq e^{-\gamma k}(f(x_0) - f^*). \tag{16}$$

We define $x_{\phi(0)} = x_0$. Assume that (16) is true at iteration $k-1$, then combining complexity bound (15) and sharpness (Sharp), for any $t \geq 1$,

$$\begin{aligned} f(x_{\phi(k-1)+t}) - f^* &\leq \frac{\kappa}{t}(f(x_{\phi(k-1)}) - f^*)^{\frac{2}{r}} \\ &\leq \frac{\kappa}{t} e^{-\gamma \frac{2}{r}(k-1)} (f(x_0) - f^*)^{\frac{2}{r}}. \end{aligned}$$

Taking $t_k = e^{\gamma(1-\tau)} \kappa (f(x_0) - f^*)^{-\tau} e^{\gamma \tau k}$ and $\phi(k) = \phi(k-1) + t_k$, (16) holds at iteration k. Using Lemma 3.1, we obtain at iteration $N = \phi(n) = \sum_{k=1}^n t_k$,

$$f(x_N) - f^* \leq \exp\left(-\gamma e^{-\gamma} \kappa^{-1} N\right)(f(x_0) - f^*), \quad \text{if } \tau = 0,$$

and

$$f(x_N) - f^* \leq \frac{f(x_0) - f^*}{(\tau \gamma e^{-\gamma} \kappa^{-1} (f(x_0) - f^*)^\tau N + 1)^{\frac{1}{\tau}}}, \quad \text{if } \tau > 0.$$

These bounds are minimal for $\gamma = 1$ and the results follow. ∎

We observe that restarting accelerated gradient methods reduces complexity from $O(1/\epsilon^\tau)$ to $O(1/\epsilon^{\tau/2})$ compared to simple gradient descent. This result is similar to the acceleration of gradient descent. We extend now this restart scheme to solve non-smooth or Hölder smooth convex optimization problem under the sharpness assumption.

4. Scheduled Restarts for Hölder smooth convex functions

In this section we seek to solve (P) assuming that the function f has a Hölder smooth gradient, i.e. satisfies (Smooth) with $s \in [1, 2]$ and $L > 0$ on a set $J \subset \mathbb{R}^n$. If the function is non smooth, it satisfies (Smooth) with $s = 1$ and L taken as the maximum norm of subgradients on J. Without further assumptions on f, the optimal algorithm to solve the convex optimization problem (P) is the Universal Fast Gradient Method in [Nesterov, 2015]. Given a target accuracy ϵ, the Universal Fast Gradient Method starts at a point x_0 and outputs after t iterations a point $y \triangleq \mathcal{U}(x_0, \epsilon, t)$, such that

$$f(y) - f^* \leq \frac{\epsilon}{2} + \frac{cL^{\frac{2}{s}} d(x_0, X^*)^2}{\epsilon^{\frac{2}{s}} t^{\frac{2q}{s}}} \frac{\epsilon}{2}, \tag{17}$$

where c is a constant ($c = 2^{\frac{4s-2}{s}}$) and

$$q \triangleq \frac{3s - 2}{2} \tag{18}$$

is the optimal rate of convergence for s-smooth functions. More details about the Universal Fast Gradient Method are given in Appendix B.1.

We will again assume that f is sharp with parameters (r, μ) on a set $K \supseteq X^*$, i.e.

$$\mu d(x, X^*)^r \leq f(x) - f^*, \quad \text{for every } x \in K.$$

As mentioned in Section 2.2, if $r > s$, smoothness or sharpness are local properties, *i.e.* either J or K or both are bounded, our analysis is therefore local. In the following we assume for simplicity, given an initial point x_0, that smoothness and sharpness are satisfied simultaneously on the sublevel set $\{x,\ f(x) \leq f(x_0)\}$. The key difference with the smooth case described in the previous section is that here we schedule *both* the target accuracy ϵ_k used by the algorithm and the number of iterations t_k made before the k^{th} restart. Our scheme is described in Algorithm H-RESTART.

Algorithm 2 General Scheduled restarts for convex minimisation **(H-RESTART)**

Inputs: $x_0 \in \mathbb{R}^n$, $\epsilon_0 \geq f(x_0) - f^*$, $\gamma \geq 0$ and a sequence t_k for $k = 1, \ldots, R$.
for $k = 1, \ldots, R$ **do**

$$\begin{array}{rcl} \epsilon_k & := & e^{-\gamma} \epsilon_{k-1} \\ x_k & := & \mathcal{U}(x_{k-1}, \epsilon_k, t_k) \end{array} \tag{H-RESTART}$$

end for
Output: $\hat{x} := x_R$

Our strategy is to choose a sequence t_k that ensures

$$f(x_k) - f^* \leq \epsilon_k,$$

for the geometrically decreasing sequence ϵ_k. The overall complexity of our method will then depend on the growth of t_k as described in Lemma 3.1.

Proposition 4.1. *Let f be a convex function satisfying (Smooth) with parameter (s, L) on a set J and (Sharp) with parameters (r, μ) on a set K. Given $x_0 \in \mathbb{R}^n$ assume that $\{x,\ f(x) \leq f(x_0)\} \subset J \cap K$. We write x_k the sequence generated by the (H-RESTART) scheme to solve (P) for given $\epsilon_0 \geq f(x_0) - f^*$, $\gamma = q$ and $t_k = C^*_{\kappa,\tau,q} e^{\tau k}$, where*

$$C^*_{\kappa,\tau,q} \triangleq e^{1-\tau} (c\kappa)^{\frac{s}{2q}} \epsilon_0^{-\frac{\tau}{q}}$$

and κ and τ are defined in (2) and (3) respectively. The precision reached at the last point $\hat{x}$ is given by,

$$f(\hat{x}) - f^* \leq \exp\left(-qe^{-1} (c\kappa)^{-\frac{s}{2q}} N\right) \epsilon_0, \quad \textit{when } \tau = 0, \tag{19}$$

while

$$f(\hat{x}) - f^* \leq \frac{\epsilon_0}{\left(\tau e^{-1}(c\kappa)^{-\frac{s}{2q}}\epsilon_0^{\frac{\tau}{q}}N + 1\right)^{\frac{q}{\tau}}}, \quad \textit{when } \tau > 0, \tag{20}$$

where $N = \sum_{k=1}^{R} t_k$ *is total number of iterations.*

Proof. Our goal is to ensure that the target accuracy is reached at each restart, *i.e.*

$$f(x_k) - f^* \leq \epsilon_k. \tag{21}$$

By assumption, (21) holds for $k = 0$. Assume that (21) is true at iteration $k-1$, combining (Sharp) with the complexity bound in (17), we get

$$\begin{aligned} f(x_k) - f^* &\leq \frac{\epsilon_k}{2} + \frac{c\kappa(f(x_{k-1}) - f^*)^{\frac{2}{r}}}{\epsilon_k^{\frac{2}{s}} t_k^{\frac{2q}{s}}} \frac{\epsilon_k}{2} \\ &\leq \frac{\epsilon_k}{2} + \frac{c\kappa}{t_k^{\frac{2q}{s}}} \frac{\epsilon_{k-1}^{\frac{2}{r}}}{\epsilon_k^{\frac{2}{s}}} \frac{\epsilon_k}{2}. \end{aligned}$$

By definition $\epsilon_k = e^{-\gamma k}\epsilon_0$ so to ensure (21) at iteration k we impose

$$\frac{c\kappa e^{\gamma\frac{2}{r}} e^{-\gamma\left(\frac{2}{r}-\frac{2}{s}\right)k}}{t_k^{\frac{2q}{s}}} \epsilon_0^{\frac{2}{r}-\frac{2}{s}} \leq 1$$

Rearranging terms in last inequality, using τ defined in (3), we get

$$t_k \geq e^{\gamma\frac{1-\tau}{q}} (c\kappa)^{\frac{s}{2q}} \epsilon_0^{-\frac{\tau}{q}} e^{\frac{\gamma\tau}{q}k}$$

Choosing $t_k = Ce^{\alpha k}$, where

$$\begin{aligned} C &= e^{\gamma\frac{1-\tau}{q}} (c\kappa)^{\frac{s}{2q}} \epsilon_0^{-\frac{\tau}{q}} \\ \alpha &= \frac{\gamma\tau}{q} \end{aligned}$$

and using Lemma 3.1 we get

$$f(\hat{x}) - f^* \leq \exp(-\gamma e^{-\frac{\gamma}{q}} (c\kappa)^{-\frac{s}{2q}} N)\epsilon_0 \quad \text{if } \tau = 0 \tag{22}$$

while

$$f(\hat{x}) - f^* \leq \frac{\epsilon_0}{\left(\frac{\gamma\tau}{q} e^{-\frac{\gamma}{q}} (c\kappa)^{-\frac{s}{2q}} \epsilon_0^{\frac{\tau}{q}} N + 1\right)^{\frac{q}{\tau}}} \quad \text{if } \tau > 0. \tag{23}$$

These bounds are minimal for $\gamma = q$ and the results follow. ■

We chose $\gamma = q$ here, however the choice of γ does not affect much bounds (22) and (23) if C and α are chosen accordingly. This matches the lower bounds for optimizing smooth and sharp functions functions [Nemirovskii and Nesterov, 1985, Page 6] up to constant factors. The rate of convergence of this method is controlled by the ratio between τ and q. If these are not known, a log-scale grid search won't be able to reach the optimal rate, even if q is known since we will miss the optimal schedule by a constant factor. If both are known, in the case of non-smooth strongly convex functions for example, a grid-search on C recovers the optimal bound. Now we will see that if f^* is known, restart produces adaptive optimal rates.

5. Restart with termination criterion

In this section we assume that we know the optimum f^* of the problem (P), or have an exact termination criterion. This is the case for example in zero-sum matrix games problems or projections on convex sets. We assume again that f is satisfies (Smooth) with parameters (s, L) on a set J and (Sharp) with parameters (r, μ) on a set K. Given an initial point x_0 we assume that smoothness and sharpness are satisfied simultaneously on the sublevel set $\{x,\ f(x) \leq f(x_0)\}$. We recall that given x_0 and a target accuracy ϵ, the Universal Fast Gradient Method outputs a point $y \triangleq \mathcal{U}(x_0, \epsilon, t)$, such that

$$f(y) - f^* \leq \frac{\epsilon}{2} + \frac{cL^{\frac{2}{s}} d(x_0, X^*)^2}{\epsilon^{\frac{2}{s}} t^{\frac{2q}{s}}} \frac{\epsilon}{2},$$

where c is a constant ($c = 2^{\frac{4s-2}{s}}$) and $q = \frac{3s-2}{2}$ is the optimal rate of convergence for s-smooth functions. As we know the optimum f^*, we can stop the algorithm when it reaches the target accuracy, *i.e.* after t_ϵ inner iterations, $y = \mathcal{U}(x_0, \epsilon, t_\epsilon)$ satisfies

$$f(y) - f^* \leq \epsilon.$$

We write $y \triangleq \mathcal{C}(x_0, \epsilon)$ the output of this method.

Here we simply restart this method and decrease the target accuracy by a constant factor after each restart. Our scheme is described in Algorithm ε-RESTART.

Algorithm 3 Restart on criterion **(ε-RESTART)**

Inputs: $x_0 \in \mathbb{R}^n, f^*, \gamma \geq 0$
Initialize: $\epsilon_0 = f(x_0) - f^*$
for $k = 1, \ldots, R$ **do**

$$\begin{array}{rcl} \epsilon_k & := & e^{-\gamma} \epsilon_{k-1} \\ x_k & := & \mathcal{C}(x_{k-1}, \epsilon_k) \end{array} \qquad (\varepsilon\text{-RESTART})$$

end for
Output: $\hat{x} := x_R$

The following result describes the convergence of this method.

Proposition 5.1. *Let f be a convex function satisfying (Smooth) with parameter (s, L) on a set J and (Sharp) with parameters (r, μ) on a set K. Given $x_0 \in \mathbb{R}^n$ assume that $\{x,\ f(x) \leq f(x_0)\} \subset J \cap K$. We write x_k the sequence generated by the (ε-RESTART) scheme to solve (P) with parameter $\gamma = q$. The precision reached at the point $\hat{x}$ is given by,*

$$f(\hat{x}) - f^* \leq \exp\left(-qe^{-1}(c\kappa)^{-\frac{s}{2q}} N\right)(f(x_0) - f^*), \quad \text{when } \tau = 0,$$

while

$$f(\hat{x}) - f^* \leq \frac{f(x_0) - f^*}{\left(\tau e^{-1}(c\kappa)^{-\frac{s}{2q}} (f(x_0) - f^*)^{\frac{\tau}{q}} N + 1\right)^{\frac{q}{\tau}}}, \quad \text{when } \tau > 0,$$

where $N = \sum_{k=1}^R t_k$ is the total number of iterations.

Proof. Given $\gamma \geq 0$, linear convergence of our scheme is ensured by our choice of target accuracies ϵ_k. It remains to compute the number of iterations t_{ϵ_k} needed by the algorithm before the k^{th} restart. Following proof of Proposition 4.1, for $k \geq 1$ we know that target accuracy is necessarily reached after

$$\bar{t}_k = e^{\gamma \frac{1-\tau}{q}} (c\kappa)^{\frac{s}{2q}} \epsilon_0^{-\frac{\tau}{q}} e^{\frac{\gamma\tau}{q} k}$$

iterations, such that $t_{\epsilon_k} \leq \bar{t}_k$. So ($\varepsilon$-RESTART) scheme achieves linear convergence while needing less inner iterates than the scheduled restart presented in Proposition 4.1. Its convergence is therefore at least as good. Here $\epsilon_0 = f(x_0) - f^*$ as we know the optimum. We choose $\gamma = q$ as it is the optimal choice derived before however bounds (22) and (23) are also valid for any choice of $\gamma \geq 0$. ■

Therefore if f^* is known, this method is adaptive, contrary to the general case in Proposition 4.1. It can even adapt to the local values of L or μ as we use a criterion instead of a set schedule. As stated in Section 4 the choice of γ does not impact much bounds (22) and (23), so running this scheme with $\gamma = 1$ makes it parameter-free while getting nearly optimal bounds.

6. Numerical Results

We test our adaptive restart method, denoted *Adap*, on several problems and compare it against simple gradient descent, accelerated gradient method, and restart heuristic enforcing monotonicity [O'Donoghue and Candes, 2015], denoted by respectively *Grad*, *Acc* and *Mono*. We plot the convergence of the best method found by grid search to compare with the restart heuristic. This implicitly assumes that the grid search is run in parallel with enough servers. Large dots represent the restart iterations. The paper focuses solely on unconstrained problem of the form (P) but our approach can directly be extended to composite problems by using the proximal variant of the gradient, accelerated and universal fast gradient methods as in [Nesterov, 2007]. This includes constrained by considering the composite problem composed by the function and the indicator function of the set.

In Figure 1, these methods are used to solve classification problems with various losses on the UCI *Sonar* and *Madelon* data sets [Asuncion and Newman, 2007] with $(n, d) = (208, 60)$ and $(2000, 500)$ respectively, the latter being poorly conditioned. For least square loss on sonar data set, we observe much faster convergence of the restart schemes compared to the accelerated method. These results were already observed by O'Donoghue and Candes [2015]. On the *Madelon* data set we observe that restart heuristics are faster but scheduled restart compare favorably while not getting stuck in a restart loop like the restart with heuristic. For logistic loss, we observe that restart does not provide much improvement. In fact for this problem gradient descent and accelerated algorithm have identical, excellent performance. For hinge loss, we regularized by a squared norm and optimize the dual. This amounts to solve a quadratic problem with box constraints. We observe here that the scheduled restart scheme offer much faster convergence, while restart heuristics may be activated too late to offer the fast convergence possible with the sharpness assumption. For the least squares penalized by a ℓ_1 norm a.k.a. the LASSO problem, we observe similar results than for the dual SVM problem which had constraints. This highlights the benefits of a sharpness assumption for these two problems and precisely quantifying this sharpness from the data structure is an interesting open problem. Regularization parameters for dual SVM and LASSO were set to one.

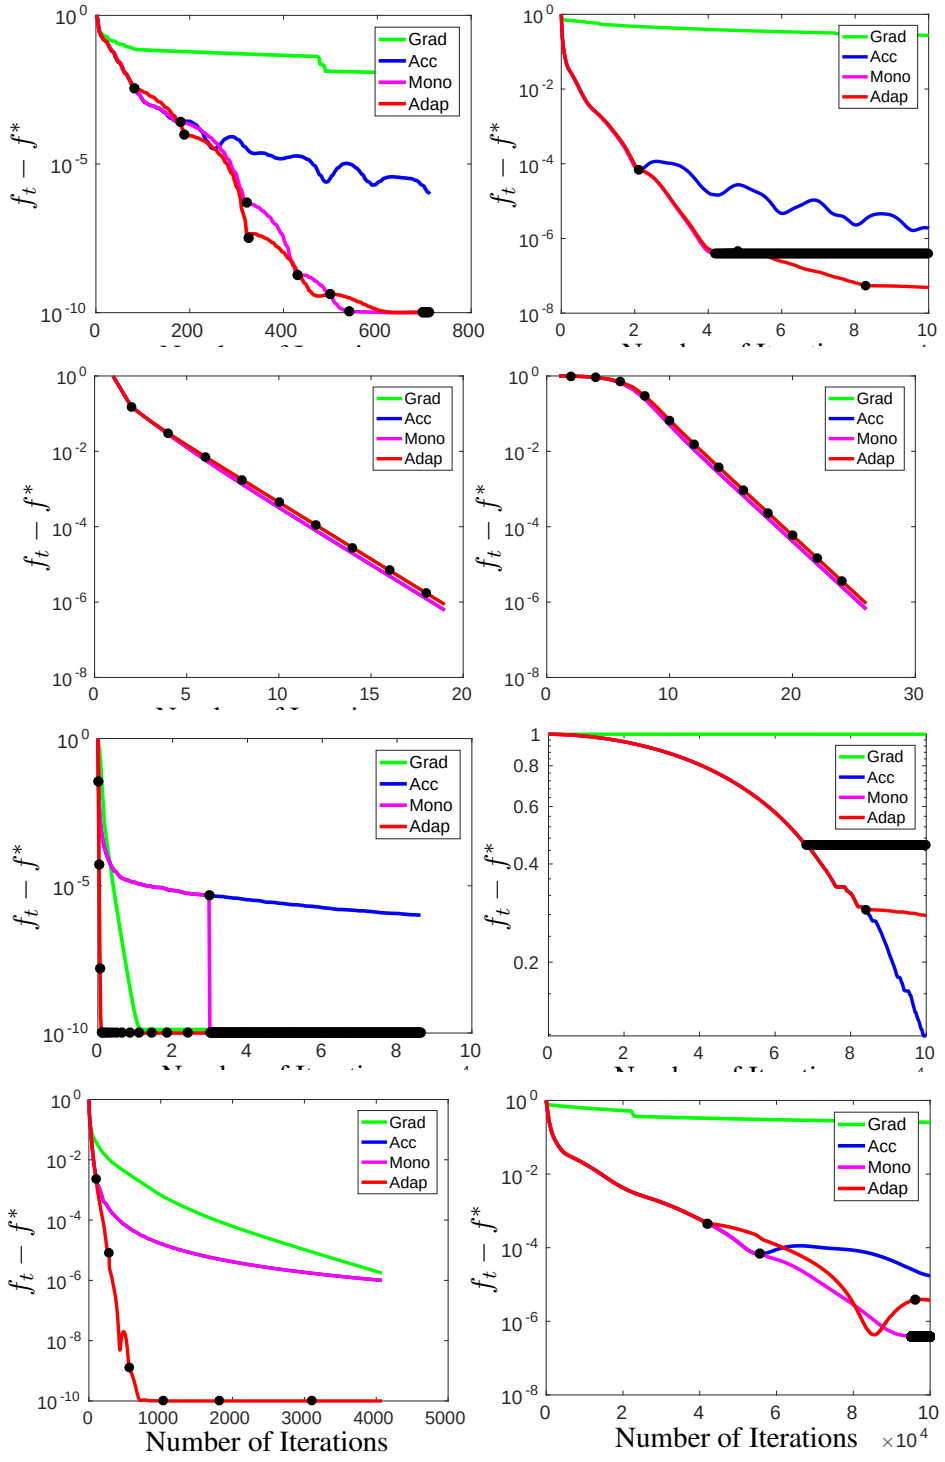

FIGURE 1. Sonar (left) and Madelon (right) data sets. From top to bottom: least square loss, logistic loss, dual SVM problem and LASSO. We use adaptive restarts (Adap), gradient descent (Grad), accelerated gradient (Acc) and restart heuristic enforcing monotonicity (Mono). Large dots represent the restart iterations

APPENDIX A. ŁOJASEVICZ'S INEQUALITY

A.1. **Łojasevicz inequality for univariate functions.** To fix ideas and illustrate the result, we present a simple proof of Łojasievicz's inequality for real-analytic functions in dimension one.

Theorem A.1. *Łojasewicz inequality [Łojasiewicz, 1963, Łojasiewicz, 1993] shows that if a function $f : \mathbf{dom}\, f \subset \mathbb{R} \to \mathbb{R}$ is analytic then for every compact set $K \subset \mathbf{dom}\, f$ there exists $r > 0$ and $\mu > 0$ such that for every $x \in K$,*

$$\mu d(x, X^*)^r \leq f(x) - f^* \tag{24}$$

Equivalently if x^* is a critical point of f, i.e. $\nabla f(x^*) = 0$, then there is $\theta \in [0, 1[$ such that

$$\frac{|f(x) - f(x^*)|^\theta}{\|\nabla f(x)\|} \tag{25}$$

is bounded in a neighborhood of x^* and compared to (24), $\theta = 1 - 1/r$.

Proof. Let us write the series expansion of f as

$$f(x) - f(y) = \sum_{k=q}^{\infty} \frac{f^{(k)}(y)}{k!}(x - y)^k$$

where $q \geq 0$ is the smallest coefficient for which $f^{(q)}(y) \neq 0$. There is an interval V around y such that

$$\frac{1}{2}\frac{f^{(q)}(y)}{q!}|x - y|^q \leq |f(x) - f(y)| \leq 2\frac{f^{(q)}(y)}{q!}|x - y|^q$$

Setting $y = x^*$ the point in X^* reaching $d(x, X^*)$, this yields inequality (24) with $r = q$. Assuming $f'(y) = 0$ (hence $q \geq 2$), and applying the bound above to the function $f'(x)$, we also get

$$\frac{1}{2}\frac{f^{(q)}(y)}{(q-1)!}|y - x|^{q-1} \leq |f'(x)|$$

which together with the inequality on $|f(x) - f(y)|$ yields

$$\frac{1}{2}\left(\frac{q!}{2}\right)^{\frac{q-1}{q}} \frac{f^{(q)}(x)^{\frac{1}{q}}}{(q-1)!}\, |f(x) - f(y)|^{1-\frac{1}{q}} \leq |f'(y)|$$

which is exactly (25) with $\theta = 1 - 1/q$ in a neighborhood V' of x. ∎

A.2. **Subanalytic functions .** We recall here the definition of a subanalytic functions and the results derived by Bolte et al. [2007]. We first define semianalytic sets.

Definition A.2. *A subset $A \subset \mathbb{R}^n$ is called semianalytic iff for each point there is a neighborhood V such that $A \cap V$ can be represented as*

$$A \cap V = \bigcup_{i=1}^{p} \bigcap_{j=1}^{p} \{x \in V : f_{ij}(x) = 0,\, g_{ij}(x) < 0\}$$

where the function f_{ij} and $g_{ij} : V \to \mathbb{R}$ are real analytic.

We then define subanalytic sets, as projections of semianalytic sets.

Definition A.3. *A subset $A \subset \mathbb{R}^n$ is called subanalytic iff for each point there is a neighborhood V such that $A \cap V$ can be represented as*

$$A \cap V = \{x \in \mathbb{R}^n : (x, y) \in B\}$$

where $B \subset \mathbb{R}^n \times \mathbb{R}^m$ is semianalytic.

Finally, a function is said to be subanalytic iff its graph is subanalytic.

Definition A.4. *A function $f : \mathbb{R}^n \to \mathbb{R} \cup \{\infty\}$ is called subanalytic if and only if its graph is a subanalytic subset of $\mathbb{R}^n \times \mathbb{R}$.*

Łojasievicz inequality is then also satisfied by lower semi-continuous subanalytic convex functions.

Theorem A.5. *Let $f : \mathbb{R}^n \to \mathbb{R} \cup \{\infty\}$ be a lower semicontinuous convex subanalytic function whose set X^* of critical points is nonempty. For any bounded set K, there is an exponent $r > 1$ and a constant $\mu > 0$ such that*

$$\mu d(x, X^*)^r \leq f(x) - f^*.$$

This last result shows that inequality (Sharp) holds for a very broad class of functions.

Appendix B. Algorithms & Complexity Bounds

B.1. **Universal Fast Gradient Method .** The Universal Fast Gradient Method [Nesterov, 2015] is detailed in Algorithm 4. Assume that f satisfies (Smooth) with parameters (s, L), then given a target accuracy ϵ, it starts at a point x_0 and outputs after t iterations a point $y \triangleq \mathcal{U}(x_0, \epsilon, t)$, such that

$$f(y) - f^* \leq \frac{\epsilon}{2} + \frac{cL^{\frac{2}{s}} d(x_0, X^*)^2}{\epsilon^{\frac{2}{s}} t^{\frac{2q}{s}}} \frac{\epsilon}{2},$$

where c is a constant ($c = 2^{\frac{4s-2}{s}}$) and

$$q \triangleq \frac{3s-2}{2}$$

is the optimal rate of convergence for s-smooth functions. The method does not need to know the smoothness parameters (s, L), but the target accuracy ϵ is used to parameterize the algorithm.

The Universal Fast Gradient Method requires an estimate L_0 of the smoothness parameter L to start a line search on L. This line search is proven to increase the complexity of the algorithm by at most a constant factor plus a logarithmic term and ensures that the overall complexity does not depend on L_0 but on L. In our restart schemes we use a first estimate L_0 when running the algorithm for the first time and we use the last estimate found by the algorithm when restarting it.

Finally if $X^* \neq \emptyset$, the Universal Fast Gradient Method produces a convergent sequence of iterates. Therefore if the Łojasievicz inequality is satisfied on a compact set K, it will be valid for all our iterates after perhaps reducing μ.

B.2. **Accelerated gradient method.** Accelerated Gradient Method is a special instance of the Universal Fast Gradient Method when the function is known to be smooth. In that case the optimal ϵ to run the Universal Fast Gradient method is 0 (otherwise it requires to know the parameters of the function). Given an initial point x_0, Accelerated Gradient Method outputs, after t iterations, a point y_t denoted $y \triangleq \mathcal{A}(x_0, t) = \mathcal{U}(x_0, 0, t)$ such that

$$f(y) - f^* \leq \frac{cL}{t^2} d(x_0, X^*)^2,$$

where $c = 4$.

B.3. **Gradient Descent.** We recall in Algorithm 5 the simple Gradient Descent when the function is smooth with constant L. It starts at point x_0 and outputs iterates $x_t = \mathcal{G}(x_0, t)$ such that

$$f(x_t) - f^* \leq \frac{L}{t} d(x_0, X^*)^2,$$

It performs a line search on the Lipschitz constant L such that L_0 can be chosen arbitrarily.

Algorithm 4 Universal Fast Gradient Method

Inputs: x_0, L_0, ϵ
Initialize: $\phi_0(x) := \frac{1}{2}\|x_0 - x\|_2^2,\ y_0 := x_0,\ A_0 := 0,\ \hat{L} := L_0$
for $t = 0, \ldots, T$ **do**
Find

$$z_t := \operatorname*{argmin}_x \phi_t(x)$$

repeat
Find $a \geq 0$, such that

$$a^2 = \frac{1}{\hat{L}}(A_t + a)$$

Choose

$$\begin{aligned} \tau &:= \frac{a}{A_t + a} \\ x &:= \tau z_t + (1 - \tau) y_t \\ y &:= x - \frac{1}{\hat{L}} \nabla f(x) \end{aligned}$$

if $f(y) \geq f(x) + \langle \nabla f(x), y - x \rangle + \frac{\hat{L}}{2}\|y - x\|_2^2 + \frac{\tau\epsilon}{2}$ **then** $\hat{L} := 2\hat{L}$ **end if**
until $f(y) \leq f(x) + \langle \nabla f(x), y - x \rangle + \frac{\hat{L}}{2}\|y - x\|_2^2 + \frac{\tau\epsilon}{2}$
Set

$$\begin{gathered} x_{t+1} := x, \quad y_{t+1} := y, \quad a_{t+1} := a, \\ A_{t+1} := A_t + a_{t+1}, \quad \hat{L} := \hat{L}/2, \\ \phi_{t+1}(x) := \phi_t(x) + a_{t+1}[f(x_{t+1}) + \langle \nabla f(x_{t+1}), x - x_{t+1} \rangle]. \end{gathered}$$

end for
Output: $y = y_T$

Algorithm 5 Gradient Descent

Inputs: x_0, L_0
Initialize: $\hat{L} := L_0$
for $t = 0, \ldots$ **do**
repeat
$x := x_t - \frac{1}{\hat{L}} \nabla f(x_t)$
if $f(x) \geq f(x_t) + \langle \nabla f(x_t), x - x_t \rangle + \frac{\hat{L}}{2}\|x - x_t\|_2^2$ **then** $\hat{L} = 2\hat{L}$ **end if**
until $f(x) \leq f(x_t) + \langle \nabla f(x_t), x - x_t \rangle + \frac{\hat{L}}{2}\|x - x_t\|_2^2$
Set

$$x_{t+1} := x, \quad \hat{L} := \hat{L}/2$$

end for

APPENDIX C. ROUNDING ISSUES

As mentioned in section 3, convergence bounds are stated for real sequences of iterate counts $(t_k)_{k=1}^{\infty}$ but in practice these are integer sequences. The following Lemma details the convergence of our schemes for an approximate choice $\tilde{t}_k = \lceil t_k \rceil$

Lemma C.1. *Suppose setting*

$$t_k = \lceil Ce^{\alpha k} \rceil, \quad k = 1, \dots, R$$

for some $C > 0$ and $\alpha \geq 0$ in the (RESTART) *scheme generates a sequence of points x_k whose objective values $f(x_k)$ converge linearly, i.e.*

$$f(x_k) - f^* \leq \nu e^{-\gamma k}, \tag{26}$$

for all $k \geq 0$ with $\nu \geq 0$ and $\gamma \geq 0$. Writing $N = \sum_{k=1}^{R} t_k$ the total number of iterations, the precision reached at the resulting point $\hat{x}$ generated by (RESTART) *is given by,*

$$f(\hat{x}) - f^* \leq \nu \exp(-\gamma N/(C+1)), \quad \textit{when } \alpha = 0,$$

and

$$f(\hat{x}) - f^* \leq \frac{\nu}{(\alpha e^{-\alpha} C^{-1} N + 1)^{\frac{\gamma}{\alpha}}}, \quad \textit{when } \alpha > 0,$$

where $N' = N - \frac{\log\left((e^{\alpha}-1)e^{-\alpha}C^{-1}N+1\right)}{\alpha}$.

Proof. At the R^{th} point generated, $N = \sum_{k=1}^{R} t_k$. If $t_k = \lceil C \rceil$, define $\epsilon = \lceil C \rceil - C$ such that $0 \leq \epsilon < 1$. Then $N = R(C + \epsilon)$, injecting it in (26) at the R^{th} point, we get

$$f(\hat{x}) - f^* \leq \nu e^{-\gamma \frac{N}{C+\epsilon}} \leq \nu e^{-\gamma \frac{N}{C+1}}.$$

On the other hand, if $t_k = \lceil Ce^{\alpha k} \rceil$, define $\epsilon_k = \lceil Ce^{\alpha k} \rceil - Ce^{\alpha k}$, such that $0 \leq \epsilon_k < 1$. On one hand

$$N \geq \sum_{k=1}^{R} Ce^{\alpha k},$$

such that

$$R \leq \frac{\log\left((e^{\alpha}-1)e^{-\alpha}C^{-1}N+1\right)}{\alpha}.$$

On the other hand,

$$\begin{aligned} N = \sum_{k=1}^{R} t_k &= \frac{Ce^{\alpha}}{e^{\alpha}-1}(e^{\alpha R}-1) + \sum_{k=1}^{R} \epsilon_k \\ &\leq \frac{Ce^{\alpha}}{e^{\alpha}-1}(e^{\alpha R}-1) + R \\ &\leq \frac{Ce^{\alpha}}{e^{\alpha}-1}(e^{\alpha R}-1) + \frac{\log\left((e^{\alpha}-1)e^{-\alpha}C^{-1}N+1\right)}{\alpha}, \end{aligned}$$

such that

$$R \geq \frac{\log\left(\alpha e^{-\alpha} C^{-1} N' + 1\right)}{\alpha}.$$

Injecting it in (26) at the R^{th} point we get the result. ■

REFERENCES

Asuncion, A. and Newman, D. [2007], 'Uci machine learning repository'.

Bierstone, E. and Milman, P. D. [1988], 'Semianalytic and subanalytic sets', *Publications Mathématiques de l'IHÉS* **67**, 5–42.

Bolte, J., Daniilidis, A. and Lewis, A. [2007], 'The lojasiewicz inequality for nonsmooth subanalytic functions with applications to subgradient dynamical systems', *SIAM Journal on Optimization* **17**(4), 1205–1223.

Burke, J., Lewis, A. and Overton, M. [2002], 'Approximating Subdifferentials by Random Sampling of Gradients', *Mathematics of Operations Research* **27**(3), 567–584.

Gilpin, A., Pena, J. and Sandholm, T. [2012], 'First-order algorithm with $\mathcal{O}(\log 1/\epsilon)$ convergence for ϵ-equilibrium in two-person zero-sum games', *Mathematical programming* **133**(1-2), 279–298.

Giselsson, P. and Boyd, S. [2014], Monotonicity and restart in fast gradient methods, *in* '53rd IEEE Conference on Decision and Control', IEEE, pp. 5058–5063.

Iouditski, A. and Nesterov, Y. [2014], 'Primal-dual subgradient methods for minimizing uniformly convex functions', *arXiv preprint arXiv:1401.1792* .

Karimi, H., Nutini, J. and Schmidt, M. [2016], Linear convergence of gradient and proximal-gradient methods under the polyak-łojasiewicz condition, *in* 'Joint European Conference on Machine Learning and Knowledge Discovery in Databases', Springer, pp. 795–811.

Kurdyka, K. [1998], On gradients of functions definable in o-minimal structures, *in* 'Annales de l'institut Fourier', Vol. 48, pp. 769–783.

Łojasiewicz, S. [1963], 'Une propriété topologique des sous-ensembles analytiques réels', *Les équations aux dérivées partielles* pp. 87–89.

Łojasiewicz, S. [1993], 'Sur la géométrie semi-et sous-analytique', *Annales de l'institut Fourier* **43**(5), 1575–1595.

Nemirovskii, A. and Nesterov, Y. E. [1985], 'Optimal methods of smooth convex minimization', *USSR Computational Mathematics and Mathematical Physics* **25**(2), 21–30.

Nesterov, Y. [1983], 'A method of solving a convex programming problem with convergence rate $O(1/k^2)$', *Soviet Mathematics Doklady* **27**(2), 372–376.

Nesterov, Y. [2003], *Introductory Lectures on Convex Optimization*, Springer.

Nesterov, Y. [2007], 'Gradient methods for minimizing composite objective function', *CORE DP2007/96* .

Nesterov, Y. [2015], 'Universal gradient methods for convex optimization problems', *Mathematical Programming* **152**(1-2), 381–404.

O'Donoghue, B. and Candes, E. [2015], 'Adaptive restart for accelerated gradient schemes', *Foundations of computational mathematics* **15**(3), 715–732.

Polyak, B. T. [1963], 'Gradient methods for minimizing functionals', *Zhurnal Vychislitel'noi Matematiki i Matematicheskoi Fiziki* **3**(4), 643–653.

Su, W., Boyd, S. and Candes, E. [2014], A differential equation for modeling nesterov's accelerated gradient method: Theory and insights, *in* 'Advances in Neural Information Processing Systems', pp. 2510–2518.

ACKNOWLEDGEMENTS

AA is at CNRS, attached to the Département d'Informatique at École Normale Supérieure in Paris, INRIA-Sierra team, PSL Research University. The authors would like to acknowledge support from a starting grant from the European Research Council (ERC project SIPA), as well as support from the chaire *Économie des nouvelles données* with the *data science* joint research initiative with the *fonds AXA pour la recherche*.

D.I., ÉCOLE NORMALE SUPÉRIEURE, PARIS, FRANCE.
E-mail address: vincent.roulet@inria.fr

CNRS & D.I., UMR 8548,
ÉCOLE NORMALE SUPÉRIEURE, PARIS, FRANCE.
E-mail address: aspremon@ens.fr